

Sep 2002

OFSD

**A Dynamic Asset Tasking
Technique for Integrated
Surveillance Operations**

Paul E. Berry

DSTO-RR-0246

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited



A Dynamic Asset Tasking Technique for Integrated Surveillance Operations

Paul E Berry

Intelligence, Surveillance and Reconnaissance Division
Information Sciences Laboratory

DSTO-RR-0246

ABSTRACT

This report is concerned with the development of a surveillance dynamic tasking tool whose purpose is to demonstrate benefits in terms of enhanced operational effectiveness through the coordinated deployment and control of a suite of surveillance assets. The underlying theory is developed for mathematically representing surveillance information, and the problem of optimising a search operation so as to maximise the information collected stated formally. A technique for solving the problem is proposed using evolutionary programming and a simplified version of the tool has been implemented so as to control sensors and platforms undertaking a simulated search operation in order to demonstrate the feasibility of the approach.

RELEASE LIMITATION

APPROVED FOR PUBLIC RELEASE

20030221 146

AQ F03-05-1018

Published by

*DSTO Information Sciences Laboratory
PO Box 1500
Edinburgh South Australia 5111 Australia*

*Telephone: (08) 8259 5555
Fax: (08) 8259 6567*

*© Commonwealth of Australia 2002
AR-012-480
September 2002*

APPROVED FOR PUBLIC RELEASE

A Dynamic Asset Tasking Technique for Integrated Surveillance Operations

Executive Summary

Integrated surveillance is concerned with the integrated management of the surveillance assets as well as the integration of the information they collect. Whereas benefits would appear *prima facie* to derive from increased integration in terms of planning, coordination and fusion, quantifying them requires a measure for the information collected in relation to the information sought. It is necessary to quantify the issue not only to *assess* improvements in surveillance effectiveness arising from increased integration, but also to *fully exploit the potential* for enhanced effectiveness, because benefits will only accrue if the potential for them is actively exploited. Improved communications and computational resources in support of surveillance operations will provide the *opportunities* for enhanced effectiveness, but it will be better planning and control decisions which *realise* the enhancements. Therefore, quantification is a prerequisite not only for the assessment of integrated surveillance but also for its full exploitation.

This document reports on progress in addressing an integrated surveillance issue in which sensors and platforms are dynamically retasked during a search operation so as to enhance their collective search effectiveness. A surveillance dynamic tasking tool has been developed and implemented for the purpose of demonstrating benefits in terms of enhanced operational effectiveness through the coordinated deployment and control of a suite of surveillance assets. The underlying theory is developed for mathematically representing surveillance information, and the problem of optimising a search operation so as to maximise the information collected is stated formally. A technique for solving the problem is proposed using evolutionary programming and a simplified version of the tool has been implemented in order to demonstrate the feasibility of the approach. An indication is given as to how the theory may be extended to deal with more general surveillance information requirements and competing prioritised requirements.

This work represents progress towards improved planning and control for integrated surveillance operations, whether civilian or military, as well as improved techniques for assessing surveillance effectiveness. It provides a scientific basis for understanding the quantitative aspects of surveillance issues. It is planned to continue this work for more general sets of surveillance information requirements and using more sophisticated optimisation techniques. It is also intended to be applied to specific surveillance problems such as Coastwatch flight route planning, UAV route optimisation and satellite surveillance operations, some of which is already in progress.

Authors

Paul E. Berry

Intelligence, Surveillance and Reconnaissance
Division

Paul studied Mathematics at the University of London, majoring in mathematical physics, before pursuing a PhD on the topic of 'Waves on Flows with Vorticity' which he completed in 1979. For 12 years he was employed at research laboratories in the UK's electricity supply industry, working on fundamental scientific problems associated with the generation and transmission of electrical energy. During this period he became progressively more involved with the modelling and optimisation of large-scale engineering systems for the purposes of investment appraisal, design evaluation, economic operation and risk mitigation. He originally joined DSTO in 1991 to work on the stochastic analysis of telecommunication networks and moved to Surveillance Systems Division in 1998 in order to assume responsibility for the assessment of integrated surveillance operations and architectures.

Contents

1. INTRODUCTION	1
2. MODELLING REQUIREMENTS.....	2
2.1 The Surveillance Information Space.....	2
2.2 The Surveillance Information Requirement	2
2.3 The Surveillance Asset Decision Space.....	2
3. DYNAMIC OPTIMISATION OF SEARCH OPERATIONS	3
3.1 Maximising information over a physical surveillance information state space	4
4. EXPERIMENTAL IMPLEMENTATION.....	8
4.1 Principles.....	8
4.2 Practice	9
5. GENERALISATION	10
5.1 Maximising information over a surveillance information space posed as a hypothesis or set of hypotheses	10
5.1.1 Example – continuous decision variable	13
5.1.2 Example – discrete decision variable.....	14
6. CONCLUSIONS.....	15
7. ACKNOWLEDGEMENTS.....	15
8. REFERENCES.....	16

1. Introduction

The purpose of this work is to demonstrate that deployed surveillance assets (sensors and their platforms) can be controlled cooperatively so as to enhance their combined effectiveness in collecting surveillance information. In order to demonstrate this an algorithm is proposed which dynamically tasks the surveillance assets based upon the information requirement and their performance in satisfying it. This requires

- A representation of the physical domain and the information collection process undertaken by the surveillance assets
- A formal representation of the surveillance information requirement and a measure for how effectively it is being achieved
- Identification of the sensor and platform decision and control variables
- The dynamic tasking algorithm which monitors past effectiveness and optimises future expected effectiveness through judicious choice of values for the decision and control variables

The dynamic tasking algorithm attempts to optimise the overall surveillance operational effectiveness resulting from sensor and platform control actions specified by decision and control variables. These will, in general, be a mix of discrete and continuous variables and are also likely to be constrained. Optimisation is required to be performed over a time horizon with the additional difficulty of targets moving unpredictably. The algorithm must perform in real time and may, in practice, never have sufficient time to generate a truly optimal solution. Therefore if the algorithm is terminated prematurely it must provide a feasible, albeit sub-optimal, solution (ie an 'Anytime Algorithm'). Also, obtaining a few good solutions is more desirable than obtaining a single perfect solution in order that human planners can choose between a range of options and consider other factors less amenable to modelling such as risk and survivability. These requirements preclude the use of standard mathematical programming techniques and so this work explores the use of biologically-inspired techniques such as evolutionary programming (EP).

It should be noted that there is a range of timescales over which decisions can be taken with regard to the deployment of sensors and their platforms. To distinguish between them the following terminology is employed: *planning* refers to the allocation of surveillance resources to a general area during a time period in advance; *scheduling* refers to the advanced specification of the times at which those resources will be deployed; *tasking* is the actual commitment of the resources to be deployed either at the scheduled times or at other times; and *control* refers to the direction in real time of sensors and platforms which are on-task. The use of the term dynamic tasking in the title refers to the tasking and control of previously scheduled assets. This implies that the assets to be used have previously been determined as well as the general area and times at which they will be deployed. The purpose of the techniques is to determine the precise timings and locations to which the assets will be deployed in response to information acquired through sensing of actual targets.

2. Modelling Requirements

2.1 The Surveillance Information Space

The *Surveillance Information Space* (*S.I.S.*) is defined as the set of all variables relevant to a particular Surveillance Information Requirement. This includes the specification of the region of interest, the time period of interest and all variables, and their ranges or allowed values, for the attributes of the targets of interest, such as total numbers of targets, their locations, speeds, tracks, types and identities. The actual values assumed by these variables will never be precisely known because of uncertainty arising from incomplete coverage, unpredictable target motion and imperfect measurements. Their values will therefore be expected to change over time as the situation evolves and as new evidence arises. The purpose of the *S.I.S.* is to represent the *process* which leads to the satisfaction of the *S.I.R.*, the aim of the present paper being to demonstrate how this process can be optimised in some sense. Note that even if a target attribute is not explicitly stated in the *S.I.R.*, it may be required for the *S.I.S.* For example, if the requirement is to identify a target, a location estimate will be needed initially to enable a suitable asset to be cued to measure this attribute.

2.2 The Surveillance Information Requirement

The purpose of surveillance is to gather information pertinent to an area and time period, but not just *all* or *any* information. Surveillance assets are a limited resource and they must therefore be directed to gather only that information which is required. The simplest way would be to specify the area and the types of targets within it, or targets moving within a particular speed range or within a range of directions. This requirement would be based upon some estimate of the threat posed by targets conforming to this requirement. The problem with this 'all or nothing' approach is that information which falls just outside such a requirement specification will be completely ignored. So for example, a target just outside the region of interest but moving towards it would be considered not to satisfy the requirement even though it will at some future point in time. A target which has been mis-classified may be deemed not to have satisfied the requirement whereas it would be better to monitor it until its classification can be confirmed. Therefore what is needed is a way of prioritising information according to its current or future relevance to the hypothesis (eg level of threat) which the surveillance operation is attempting to confirm or deny.

2.3 The Surveillance Asset Decision Space

Whereas the Surveillance Information Space is composed of dependant variables which are a function of the operational situation and sensor performances and whose values ultimately contribute to the satisfaction of the *S.I.R.*, the Surveillance Asset Decision Space *D* is composed of independent control variables and their allowable values. The decision variable values are set at discrete instants of time or are controlled

continuously over time, and are determined so as to best satisfy the *S.I.R.* (ie achieve optimality) subject to operational, environmental and resource constraints (ie maintain feasibility) over a given time horizon. In order to gauge how well a *S.I.R.* is being met and plan to satisfy it optimally in the future, a Measure of Surveillance Effectiveness (*M.S.E.*) is required. As will be seen, a *M.S.E.* differs from most military Measure of Effectiveness in needing to be a measure of information, specifically information derived from sensors. If there is prior knowledge about the likely existence of targets and their location, for example from a normalcy database, then this should be taken into account when deciding where to deploy surveillance effort since it is obviously better to look where targets are expected to be than where they are not.

3. Dynamic Optimisation of Search Operations

The surveillance tasking tool in its current state of implementation deals with the first stage of surveillance, namely search, where the objective is to collect as much information as possible, consistent with the *S.I.R.*, about the existence of targets within a region during a time period, which may subsequently be followed up by additional surveillance activities. For the search mode of a surveillance operation the tool must automatically account for the need to minimise overlap between sensors, unless overlap is required to provide the necessary level of confidence in information obtained, and not to search areas which have already been previously searched, unless required by the fact that previously undetected targets have drifted into the area from elsewhere since last searched.

It is assumed that the region of interest \mathcal{R} is divided into cells labelled $i=1,\dots,N$. The shape of the region is immaterial but the cells can be considered to be squares of uniform size. Note the cells are taken to be sufficiently small that there can be no more than one target in each cell, or at least the probability of such an occurrence is exceedingly small and to all practical purposes can be ignored. The cell size may be related to the pixel size in an image. Each cell has a probability associated with it of containing a single target p_i^n at time epoch n . This probability distribution for the existence and location of targets is initialised to the *a priori* distribution p_i^0 at time $t=0$. As sensors sweep, scan and dwell upon cells within \mathcal{R} , the probabilities for the occupation of the cells by targets change depending upon whether detections occur or not. If it is assumed for present purposes that sensors are perfect, then a detection at epoch n in cell i will set $p_i^n=1$, whereas a non-detection will set it to 0. By a perfect sensor is meant that a detection occurs in a cell if and only if there is a target in that cell, assuming that the sensor 'looks' in that cell. If a cell is not inspected then the probability of a target residing within it is assumed not to change. Note that these assumptions can be generalised. One can update probabilities based upon imperfect sensors using Bayes' Rule and can modify target probability distributions to take

account of target motion using a Markovian target motion model (eg Gauss-Markov). These issues are described in detail in Berry *et al* [1].

A surveillance asset tasking and control algorithm such as that described in [1] and the present paper could be useful for actual surveillance planning, tasking and control operations, or could be used to task and control surveillance assets within operational simulations for the purpose of assessing surveillance effectiveness. A coarser, but more efficient, analysis approach for undertaking the latter is the *Integrated Surveillance Assessment Model* described in Berry [2].

In order to provide a unified framework for the development and application of a tasking tool, the problem is approached from a probabilistic perspective as probability theory is able to address all of the modelling issues referred to above in a theoretically consistent fashion. These ideas are developed in the following sections.

3.1 Maximising information over a physical surveillance information state space

The objective of a search operation is to gather information about the existence, and non-existence, of targets within a given area. To gauge how well this objective is being satisfied and to satisfy it in the shortest time or with the minimum of resources it is necessary to quantify it. To this end, a measure for the information derived from the sensors undertaking the search operation will be defined, namely entropy. From the probabilities for the occupation of single cells comprising \mathfrak{R} a single state space S is constructed for the numbers of targets and their locations and a probability distribution defined over S :

$$\Pr\{s_{j_1, j_2, \dots, j_k}\} = \Pr\{\text{target in cells } j_1, j_2, \dots, j_k\} = \prod_{l=1}^k p_{j_l} \prod_{\substack{i=1 \\ i \neq j_l}}^N (1 - p_i)$$

This is true if it is assumed that all cells are independent in the sense that information about the existence of a target in one cell has no effect on the likelihood of targets in any other cells. This would not be correct if targets were correlated with each other, for example ships comprising a naval fleet, or if the total number of targets were assumed known. This is referred to as the *local* formulation of the target location problem and leads to some simplifying features (see Berry *et al* [1]). It is, in fact, a two dimensional Bernoulli distribution which, in the continuous limit, becomes a two dimensional Poisson distribution.

The sum of these probabilities over S is, of course, unity:

$$\sum_{s_{j_1, j_2, \dots, j_k} \in S} \Pr\{s_{j_1, j_2, \dots, j_k}\} = 1$$

and the entropy corresponding to a particular probability distribution is:

$$h(\underline{p}_S) = - \sum_{s_{j_1, j_2, \dots, j_k} \in S} \Pr\{s_{j_1, j_2, \dots, j_k}\} \ln \Pr\{s_{j_1, j_2, \dots, j_k}\}$$

where the summation is over all possible states of S and p_s is the vector of all probabilities over S . Substituting for $\Pr\{s_{j_1, j_2, \dots, j_k}\}$ one obtains, in the case of independent cells,

$$h(\underline{p}_S) = -\sum_{j=1}^N [p_j \ln p_j + (1 - p_j) \ln(1 - p_j)]$$

which is simply the sum of the entropies for the individual cells, each regarded as a state space in its own right. This expression is readily computed; furthermore, the linearity in contributions to total entropy from individual cells means that it is easy to specify entropy for sub-regions $\mathcal{R}' \subset \mathcal{R}$ if only targets occupying \mathcal{R}' are of interest. This expression provides a quantitative measure of surveillance effectiveness. It will be seen that using this expression it is also relatively easy to predict the expected entropy arising from a search decision.

As observations of targets are made the cell probabilities are updated using Bayes' rule, but not the cells where no observations occur, because cells are assumed independent. Because targets are likely to move between observations, the probabilities for their locations should be changed to reflect this according to appropriate target motion models (see Berry *et al* [1]).

This probability distribution is a representation of the search *S.I.R.* and the entropy is a measure of how well it has been satisfied in terms of either confirming or denying the existence of targets in \mathcal{R} and their locations. If the objective of a surveillance search operation is to make the information about targets and their locations in \mathcal{R} as complete as possible then the value of the entropy must be minimised. If it ever achieves the value of zero then the information obtained will be complete and lacking in any uncertainty. In practice this will never be achieved because of imperfect sensors, unpredictable target motions and inadequate coverage so the objective must be to minimise it as far as is practically possible. Any choice of sensor action d in the decision space D must be such as to minimise the expected entropy which will result from the action.

Let $p_s^n = \Pr\{s \in S \text{ at time epoch } n\}$ for an element of S , and $\underline{p}_S^n = \{p_s^n\}_{s \in S}$ be the vector of all such probabilities over S at time epoch n . For a choice of decision $d \in D$ there is a set of measurements or observations $M(d)$ possible, each element $m \in M(d)$ of which has probability $P(m, d)$ of occurring. For each decision $d \in D$ and consequent measurement set $m \in M(d)$, the probabilities are updated using Bayes' Rule

$$p'_s = \frac{L(m|s, d) p_s}{P(m, d)}$$

where the likelihood function $L(m|s, d)$ is defined as the probability that measurement set $m \in M(d)$ will occur given that the system is in state $s \in S$ and decision $d \in D$ is taken. Depending upon how the state space is defined, s could be simply the existence of a target within a cell, or a pattern of target distribution within

\mathfrak{R} . Note that the probability of a particular measurement set m occurring can be obtained from the likelihood function thus

$$P(m, d) = \sum_{s \in S} L(m | s, d) p_s$$

The entropy of the system following the measurement set m which, in turn, depends upon the decision d , is given by

$$h(\underline{p}'_s | m, d) = - \sum_{s \in S} p'_s \ln p'_s$$

This can be used to predict what the entropy will be in advance of an observation for a choice of decision d by regarding the entropy function as a random variable over the measurement space $M(d)$. Then its expected value is

$$h_e(\underline{p}'_s | d) = E_{m \in M(d)} h(\underline{p}'_s | m, d) = \sum_{m \in M(d)} P(m, d) h(\underline{p}'_s | m, d)$$

If the preferred action $d' \in D$ is such as to minimise the expected entropy at the next epoch, then it satisfies

$$h_e(\underline{p}'_s | d') = \min_{d \in D} h_e(\underline{p}'_s | d)$$

or, equivalently

$$E_{m \in M(d')} h(\underline{p}'_s | m, d') = \min_{d \in D} E_{m \in M(d)} h(\underline{p}'_s | m, d).$$

This analysis has been developed for a completely general state space and set of observations. In practice the size of the state space and the set of all possible observations may be combinatorially large which could render the computation of the expected entropy difficult, particularly so because the expected entropy would to be evaluated for a potentially large number of possible decisions, enabling the optimal one to be selected. However, a general expression may be derived for the expected entropy which simplifies its evaluation somewhat, and also enables approximations to it to be derived systematically.

Taking the expression for the entropy following an observation

$$h(\underline{p}'_s | m, d) = - \sum_{s \in S} p'_s \ln p'_s$$

and substituting for p'_s using Bayes' Rule

$$p'_s = \frac{L(m | s, d) p_s}{P(m, d)},$$

the following expression is obtained

$$h(\underline{p}'_s | m, d) = - \frac{1}{P(m, d)} \sum_{s \in S} \{ p_s L(m | s, d) \ln L(m | s, d) + L(m | s, d) p_s \ln p_s \} + \ln P(m, d)$$

Substituting this into the expression for the expected entropy

$$h_e(\underline{p}'_s | d) = \sum_{m \in M(d)} P(m, d) h(\underline{p}'_s | m, d)$$

and simplifying, yields

$$h_e(\underline{p}'_s | d) = h(\underline{p}_s) - \sum_{s \in S} p_s \sum_{m \in M(d)} L(m | s, d) \ln L(m | s, d) + \sum_{m \in M(d)} P(m, d) \ln P(m, d)$$

where $h(\underline{p}_s)$ is the entropy of the system immediately prior to the observations based upon the *a priori* probabilities. Note that this expression relates the expected *a posteriori* entropy to the *a priori* entropy and therefore expresses the difference between them which is to be maximised in order to maximise the improvement in information.

In the special case of independent cells which is proposed as the state space for the search problem, the expected entropy can be computed for each cell independently and then summed over all of the cells. The fact that this is possible was pointed out previously.

For cell i there are two possible states s (the cell is either occupied by a target or it is not with probabilities p_i and $1 - p_i$ respectively) and two possible measurements m (either a target is observed or it is not). The likelihood function $L(m, d)$ is therefore

$$L(\text{target observed} | \text{target exists}) = p_d$$

$$L(\text{target not observed} | \text{target exists}) = 1 - p_d$$

$$L(\text{target observed} | \text{target does not exist}) = p_{fa}$$

$$L(\text{target not observed} | \text{target does not exist}) = 1 - p_{fa}$$

The contribution to the expected entropy for each cell inspected (no Bayesian updates occur for cells not inspected due to the independence assumption) is therefore

$$\begin{aligned} h_{e,i} = & h_i - p_i [p_d \ln p_d + (1 - p_d) \ln(1 - p_d)] \\ & - (1 - p_i) [p_{fa} \ln p_{fa} + (1 - p_{fa}) \ln(1 - p_{fa})] \\ & + [p_d p_i + p_{fa} (1 - p_i)] \ln [p_d p_i + p_{fa} (1 - p_i)] \\ & + [1 - p_d p_i - p_{fa} (1 - p_i)] \ln [1 - p_d p_i - p_{fa} (1 - p_i)] \end{aligned}$$

This assumes that information regarding all targets is required everywhere in \mathcal{R} . In practice the requirement may be for information about targets within a strict subregion $\mathcal{R}' \subset \mathcal{R}$ and that sensor actions should therefore be optimised within this subregion. Of course information may happen to be obtained about targets outside of \mathcal{R}' but still within \mathcal{R} and which may be relevant to some future surveillance information requirement. Let S' consist of the states in S corresponding to targets in \mathcal{R}' , that is $s_{j_1, \dots, j_k} \in S'$ if and only if at least one of cells j_1, \dots, j_k is in \mathcal{R}' . Then probabilities for occupation of states in S can be conditioned on S' as follows:

$$p(s | s \in S') = \frac{p(s)}{p_{S'}}$$

$$\text{where } p_{S'} = \sum_{s \in S' \neq \emptyset} p_s,$$

that is the probabilities are summed over the states in S which contain s' .

As will be seen in what follows it is possible to define an information requirement more precisely than simply all target information relevant to a region or subregion.

4. Experimental Implementation

4.1 Principles

The analysis demonstrates that it is possible to mathematically define an objective function which represents the quantity of the information obtained about targets within a region. Furthermore it is possible, given the imperfections of sensors, to task the sensors in advance so to observe at those locations where the expected improvement in information is likely to be the greatest. It is also possible, in doing so, to take account of the physical constraints placed upon where sensors may observe by virtue of the platforms they reside upon. This could, in principle, be applied to the forward planning of surveillance flight routes, for instance, but in the present paper interest is confined to the problem of dynamically reassigning airborne sensors whose routes have already been scheduled.

In formulating this problem meaningfully it is necessary to consider what changes would be possible to the tasking of a surveillance asset and what benefits would be derived from such changes. If an optimal route has already been scheduled to meet a requirement then why would it be necessary to change it? This can only be because the assumptions underlying the original plan have lost their validity. This could be because of unforeseen delays in platform take-off times, deviations in platform speeds due to headwinds, unexpected localised environmental conditions affecting sensor performances, additional information provided by opportunistic assets such as multi-role platforms, and changing information requirements or priorities. Alternatively, the original search plan may be sub-optimal and therefore capable of improvement.

Given that changes may be required to previously planned surveillance flight routes, the next consideration is to the choice of times at which such changes should be made and how far ahead in time the changes should be planned for, namely the choice of planning horizon. It is assumed that changes to routes are implemented by moving the next defined waypoint for a platform and that all remaining future waypoints remain unchanged. The amount of the deviation can be constrained. This avoids the problem of replanning the entire route and maintaining the constraint that a platform must return to an airfield before exhausting its fuel load, which would be computationally expensive in real time. This has the effect of constraining a platform not to deviate too far from its original planned route.

It is possible, in principle, then, to generate a well-defined dynamic tasking problem in terms of an objective function and a set of constraints which aims to maximise the information collected. Its implementation within a simulated surveillance operation can be expected to demonstrate benefits in dynamically tasking sensors and platforms.

However, a suitable optimisation algorithm is required which although it may not provide truly optimal solutions, is at least capable of yielding good quality solutions in real time.

The above methodology has been applied to a specific surveillance asset tasking problem, namely optimising the allocation resources provided by a constellation of surveillance satellites for the purpose of tracking surface maritime targets (Berry & Fogg [6]). This problem is characterised by complexity in satellite orbit computation but relative simplicity in terms of the planning horizon because one only needs to plan ahead to the next epoch at which a satellite flies past the region of interest.

4.2 Practice

Prior to the development of the above analysis, an experimental tasking tool was developed within SSD to gain familiarity with the application of existing methodologies used to address this type of problem, and to explore the required structure to enable connectivity to ISAT (Integrated Surveillance Assessment Tool) simulations. Algorithm options included Genetic Algorithms, Neural Networks and Evolutionary Programming (EP) [3]. The latter technique, the simplest, was chosen for this exploratory work.

Genetic methods store multiple solutions to a problem each solution referred to as member of a population. Associated with each member is its *fitness*, which is simply a measure of how well this solution solves the problem at hand. Throughout the search for the optimal solution, a 'survival of the fittest' procedure is used, meaning that a solution with a higher fitness is chosen over one with a lower fitness. The main difference between a Genetic Algorithm and Evolutionary Programming is how new solutions (or *offspring*) are generated from existing members. In the former, two solutions are *mated* to form a new solution, while in the latter, each member of the population generates a offspring by *mutation*.

The current problem consisted of determining how best to task a given set of mobile platform-mounted sensors to achieve a specified set of surveillance objectives, over a given area, against a specified set of targets. The scenario was based on a version of the test case described in [4], but extended to allow more detailed discussion of surveillance issues.

The technique consists simply of starting with an initial tasking of assets in terms of prescribed start times and waypoints, and then perturbing and measuring effectiveness at regular intervals of time. After each perturbation, for each sensor system, if there is a gain that is greater than a random number in the range $[-p, 0]$ the alteration is accepted. In this implementation, p was chosen as 5%. Note that allowing negative gain facilitates escaping from local minima.

The EP algorithm was implemented using MATLAB and required a single input structure containing asset state information such as type, positional information, waypoints, etc. The output is also a single structure containing asset information, but only includes future waypoint details.

In addition, software was developed to enable sensor requirements and achievements to be displayed as two-dimensional colour displays. The difference between required and achieved surveillance could also be displayed. This software was used to demonstrate the potential and ease of use of such aids to tasking surveillance assets. Examples using this process showed consistent improvement towards achieving the desired surveillance capability.

As part of this work JAVA-to-MATLAB middleware was developed to enable connectivity between a JACK intelligent software agent and the tasking tool function. This allows the future use of such tools in STAGE simulations (STAGE is a high fidelity operational simulation model) (see Clark [5]).

5. Generalisation

This section has been included in order to indicate how the surveillance information requirement can be generalised beyond simply the existence and location of targets within a defined region.

5.1 Maximising information over a surveillance information space posed as a hypothesis or set of hypotheses

Consider first the representation of the *S.I.R.* This is assumed to be stated for the purpose of supporting or denying a hypothesis which is of significance to the end user, such as a threat hypothesis, and therefore embodies the end user's preference for information. The hypothesis can either be true or false and the purpose of undertaking the surveillance operation is to determine which is correct. If a hypothesis of significance is true then some appropriate response will be required, but if not then it can be safely ignored. The intention is that the decision point at which action is required to be taken should be achieved as quickly as possible within the resource limitations and surveillance asset capabilities so as to enable the operational response to be as effective as possible.

In practice there are likely to be a number of competing requirements for surveillance information so these are distinguished between by use of the subscript r . For present purposes it will not be necessary to understand the meaning of such hypotheses, only their quantitative relationship to the *S.I.S.* variables. The simplest way of developing this idea is to assume that certain cells being occupied by targets are fully consistent with a hypothesis H and remaining cells, whether occupied or not, are of no relevance to it. Note that each cell has its own probability of occupation which is unrelated to the

probability of occupation of any other cell. This is equivalent to assuming that all targets are independent. Then one may assign a random variable I_i^n to each cell i which is either 0 if the cell is unoccupied or 1 if it is. Then the probability of it being occupied at time epoch n is $p_i^n = \Pr\{I_i^n = 1\}$. The global S.I.S. is obtained by compounding together all of the individual cell states thus:

$$S = I_1 \otimes I_2 \otimes \dots \otimes I_N$$

so that if targets are truly independent the probability of the occupation of a global state can be obtained from the product of the occupation of the appropriate local states. The hypothesis H may be considered to be a subset $S_H \in S$ so that if a given $s \in S_H \subset S$ then H is true and if $s \in S_H^c \subset S$ then H is false, where $S_H \cup S_H^c = S$ of course. This may be expressed formally as

$$\Pr\{H \mid s \in S_H\} = 1$$

$$\Pr\{H \mid s \in S_H^c\} = 0$$

so that at any time epoch n the truth or otherwise of the hypothesis H can be estimated thus

$$\Pr\{H\} = \sum_{s^n \in S} \Pr\{H \cap s^n\} = \sum_{s^n \in S} \Pr\{H \mid s^n\} \Pr\{s^n\} = \sum_{s^n \in S_H} \Pr\{s^n\}$$

This is the 'all or nothing' approach referred to previously. The occupation of certain states fully supports the hypothesis and occupation of any of the rest deny it.

Other examples of hypotheses H are the existence of a specified number k of targets within \mathfrak{R} , or the existence of at least one target in \mathfrak{R} . Let H_k be the hypothesis that exactly k cells are occupied, and S_k be the space corresponding to any k cells being occupied and the remainder unoccupied. Then

$$\Pr\{H_k\} = \sum \Pr\{s \in S_k\}$$

The probability that at least one cell is occupied is then

$$\Pr\{H_{>0}\} = \sum_{k \neq 0} \Pr\{H_k\} = 1 - \Pr\{H_0\}$$

In general, one can compute the entropy of a hypothesis H regarded as a random variable to determine how strongly the evidence supports or denies its truth. If it is determined that it is likely to be true, or that the risk or ignoring its truth is operationally too high, one can also condition the state of the S.I.S. on its truth thus:

$$\Pr\{s^n \mid H\} = \frac{\Pr\{H \mid s^n\} \Pr\{s^n\}}{\sum_{s'} \Pr\{H \mid s'\} \Pr\{s'\}}$$

This is simply conditioning the probability distribution on a subset S_H of S . The purpose in doing this is to enable future sensor actions to be optimised so as to either confirm or deny the truth of a hypothesis regarding the spatial distribution of a collection of targets, rather than just their existence and individual locations. It should be noted that the probability distribution for the target locations conditioned upon the

hypothesis H can be updated simply by updating the probabilities for the individual cells and then conditioning.

More generally one would wish to associate a hypothesis more strongly with some states than others but not to some at the total exclusion of the rest. Therefore define H as a random variable which is either true or false and is correlated with the states of S . This is done by creating the joint state space $\{T, F\} \otimes S$ and then specifying the probabilities on it, namely

$$\Pr\{H = T, s\} \text{ and } \Pr\{H = F, s\}$$

where

$$\sum_{s \in S} [\Pr\{H = T, s\} + \Pr\{H = F, s\}] = 1$$

From these correlations between H and S , which can be based upon historical data, one can compute the probability of the hypothesis H being true for a given state $s \in S$ being assumed to be true:

$$\Pr\{H = T | s\} = \frac{\Pr\{H = T, s\}}{\Pr\{H = T, s\} + \Pr\{H = F, s\}}$$

Note that this is a generalisation of the above where the conditional probabilities were either 0 or 1. From these conditional probabilities one can compute at any time epoch n the probability that the hypothesis H is true:

$$\Pr\{H = T\} = \sum_{s \in S} \Pr\{H = T | s^n\} \Pr\{s^n\}$$

and, as before, condition the probabilities over the state space on the assumption that the hypothesis is true:

$$\Pr\{s^n | H = T\} = \frac{\Pr\{H = T | s^n\} \Pr\{s^n\}}{\Pr\{H = T\}}$$

If it is decided to deploy sensors so as to collect information for the purpose of confirming or denying the hypothesis H then the overall objective must be to minimise the entropy associated with H , namely $h(H)$ defined by:

$$h(H) = -\Pr\{H = T\} \ln \Pr\{H = T\} - \Pr\{H = F\} \ln \Pr\{H = F\}$$

Hence at any time epoch n the actions taken should be such as to cause the maximum incremental reduction in the expected entropy following the target motions and sensor actions:

$$\max_{d \in D} E_{e(d)} [h(H^n) - h(H^{n+1})]$$

where D is the decision space, d is an element of D and $e(d)$ is evidence arising from sensor decision d . From any given sensor decision d the evidence e , ie set of measurements, arising from the sensors is stochastic in nature and hence the need for an expectation over all possible evidence. For each possible set of values for e the state s^{n+1} is computed using Bayesian inference. If it is assumed that the sensors are perfect then the expectation and Bayesian inference is rendered redundant because the measuring process is deterministic, that is the outcome of a measurement is that a target exists within an inspected cell with probability 0 or 1 irrespective of its prior probability. Although the problem has been simplified through this assumption a

probabilistic approach is still necessary because probabilities still lie between 0 and 1 where sensors have not yet searched and because non-deterministic target motion has the effect of spreading out distributions.

Intuitively one would expect to have sensors search areas not previously searched or at least those areas not previously searched which are most relevant to the truth or otherwise of the hypothesis. If the hypothesis embodies the assumption that there is a given number of targets, then confirming that they are not in an easily accessible area will help to prove that they are elsewhere in a location which may be less easily accessible (referred to as negative information).

5.1.1 Example – continuous decision variable

One can see how this idea works if one assumes a single continuously-valued decision variable d which is constrained to make small changes between epochs. This is a simplified example chosen for the purpose of demonstration. In reality there will be several decision variables and some will be constrained to assume discrete values. The intention is to maximise $h(H^n) - h(H^{n+1}) = h(H(s(d^n))) - h(H(s(d^{n+1})))$ over d^{n+1} . Performing a Taylor series expansion about d^n we have

$$h(H(s(d^{n+1}))) = h(H(s(d^n))) + (d^{n+1} - d^n) \frac{\partial h(H(s(d)))}{\partial d} \Big|_{d=d^n}$$

to first order. Hence

$$\begin{aligned} & h(H^n) - h(H^{n+1}) \\ &= -(d^{n+1} - d^n) \frac{\partial h(H(s(d)))}{\partial d} \Big|_{d=d^n} \\ &= -(d^{n+1} - d^n) \frac{\partial}{\partial d} [-\Pr\{H = T\} \ln \Pr\{H = T\} - \Pr\{H = F\} \ln \Pr\{H = F\}] \Big|_{d=d^n} \\ &= (d^{n+1} - d^n) \left[(1 + \ln \Pr\{H = T\}) \frac{\partial \Pr\{H = T\}}{\partial d} + (1 + \ln \Pr\{H = F\}) \frac{\partial \Pr\{H = F\}}{\partial d} \right] \\ &= (d^{n+1} - d^n) \ln \left[\frac{\Pr\{H = T\}}{1 - \Pr\{H = T\}} \right] \frac{\partial \Pr\{H = T\}}{\partial d} \end{aligned}$$

The intention is to maximise this expression w.r.t. d^{n+1} subject to d^{n+1} being constrained. In fact d^{n+1} is constrained to be close to d^n in order that this expression be valid as it has resulted from the truncation of a Taylor series.

Now $\ln \left[\frac{\Pr\{H = T\}}{1 - \Pr\{H = T\}} \right]$ is 0 when $\Pr\{H = T\} = \Pr\{H = F\} = 0.5$, is positive when

$\Pr\{H = T\} > 0.5$ and negative when $\Pr\{H = T\} < 0.5$. This says that if H is more likely to be true than false then d should be changed in a direction which improves the probability that H is true and, conversely, if H is less likely to be true than false then d

should be changed so as to reduce the probability that H is true. If $\Pr\{H = T\} = \Pr\{H = F\} = 0.5$ then the second order term in the Taylor series needs to be inspected.

5.1.2 Example – discrete decision variable

In practice $\Pr\{H=T\}$ will depend upon the variables comprising the decision space in complex ways so for present purposes it will continue to be assumed that the sensors are perfect. Suppose that the hypothesis H is the statement that there is at least one target in \mathfrak{R} so that if $H=F$ there are no targets in \mathfrak{R} . Then

$$\Pr\{H = T\} = 1 - \prod_i (1 - p_i)$$

Note that only one target needs to be detected in order for the hypothesis to be deemed true. If the decision d is the label for the next cell to be inspected and only one cell can be inspected at each epoch, then intuitively the best option is to inspect a cell which has the highest probability of containing a target because when $p_{i=d} = 1$, $\Pr\{H = T\} = 1$. An alternative strategy is to attempt to prove that $\Pr\{H = T\} = 0$ but this involves inspecting every cell in \mathfrak{R} to ensure that there is no target in it. This will eventuate anyway if in the course of trying to prove that $\Pr\{H = T\} = 1$ no cells are found to contain targets.

This can be formalised as follows. Note that when a cell is inspected and no target is observed then its probability of occupation becomes zero. Irrespective of the value of $\Pr\{H = T\}$ before the inspection, the expected consequence of inspecting cell $i=d$ is $\Pr\{H = T\} = p_{i=d}$, the probability that there is a target in the cell inspected. If a target is detected in cell $i=d$ then

$$\Pr\{H = T \mid \text{detection in } i = d\} = 1$$

but if none is detected then

$$\Pr\{H = T \mid \text{no detection in } i = d\} = 1 - \prod_{i \neq d} (1 - p_i)$$

Since the entropy for hypothesis H is, in general:

$$h(H(s)) = -\Pr\{H = T\} \ln \Pr\{H = T\} - [1 - \Pr\{H = T\}] \ln [1 - \Pr\{H = T\}],$$

the entropy following a detection in cell $i=d$ is:

$$h(H \mid \text{target observed in } i = d) = 0$$

and the entropy following no detection in cell $i=d$ is:

$$\begin{aligned} h(H \mid \text{no target observed in } i = d) = \\ -[1 - \prod_{i \neq d} (1 - p_i)] \ln [1 - \prod_{i \neq d} (1 - p_i)] - \prod_{i \neq d} (1 - p_i) \ln \prod_{i \neq d} (1 - p_i) \end{aligned}$$

Hence the expected entropy prior to observing cell $i=d$ is:

$$\begin{aligned}
E\{h(H)\} &= p_{i=d}h(H \mid \text{detection in } d) + (1 - p_{i=d})h(H \mid \text{no detection in } d) \\
&= -(1 - p_{i=d})\left\{ \left[1 - \prod_{i \neq d} (1 - p_i)\right] \ln \left[1 - \prod_{i \neq d} (1 - p_i)\right] + \prod_{i \neq d} (1 - p_i) \ln \prod_{i \neq d} (1 - p_i) \right\}
\end{aligned}$$

If we put $1 - p_{i=d} = \alpha$ and $\prod (1 - p_i) = \pi$, where the product is over all i , then this simplifies to

$$E\{h(H)\} = -\alpha \ln(1 - \pi / \alpha) + \pi \ln(\alpha / \pi - 1)$$

The decision to be made is the choice of cell d to inspect which will maximise the expected entropy. Note that the parameter π is fixed and independent of the choice of cell whereas α does depend upon the choice of cell through its probability. This could be solved by computing the expected entropy for all cells and selecting that cell which has the least value. However it will be shown that the expected entropy is a monotonically increasing function of α and hence decreases with increasing $p_{i=d}$. Hence the choice of cell should be that which has the highest probability as originally proposed. By differentiating the expression for the expected entropy

$$\frac{\partial E\{h(H)\}}{\partial \alpha} = -\ln(1 - \pi / \alpha) > 0,$$

it is seen that the function is a monotonically increasing function of α .

6. Conclusions

It has been demonstrated that it is possible to specify the state of a surveillance search operation probabilistically, evolve the state space probabilities over time according to the likely target behaviours, update the state space probabilities in the light of evidence from imperfect sensors, and define an objective function of the state space probabilities in terms of entropy which quantitatively expresses the level of completeness and accuracy of the target information. By scheduling sensors so as to minimise the expected entropy it is possible, in principle, to maximise the quality of the information obtained. An evolutionary programming technique has been implemented in conjunction with a simulated surveillance scenario to demonstrate the feasibility of the approach, albeit for an heuristically defined objective function. The specification of the sensor scheduling problem has been extended to indicate how more general surveillance information requirements can be accommodated within this formulation. Although the feasibility of the approach has been established, further work is required to fully implement it and demonstrate its benefits for scenarios of practical interest.

7. Acknowledgements

Thanks are due to David Fogg for his encouragement of this work and his suggestion that entropy would be an appropriate measure of surveillance information. The exploratory experimental work described in Section 4.2 was carried out by Victor Fok.

8. References

- [1] Berry, P.E., Fogg, D. and Pontecorvo, C. "GAMBIT: Gauss-Markov and Bayesian Inference Technique for Information Uncertainty and Decision-Making in Surveillance Simulations", DSTO Research Report in draft
- [2] Berry, P.E. "Integrated Surveillance Model for the Assessment of Surveillance Architectures and Operations", DSTO Research Report in draft
- [3] Porto, V.W. "Using Evolutionary Programming to Optimize the Allocation of Surveillance Assets", Second Asia-Pacific Conf on Simulated Evolution and Learning, Canberra, November 1998, Selected Papers (Publ: Springer, Eds McKay et alia)
- [4] Hall, D & Berry, P E "Merits of Monte Carlo Simulation for the AIR6000 Test Scenario", SSD Divisional Discussion Paper, 1999
- [5] Clark, D "Sea-Air Gap Agent" DSTO-TR-1304, 2002
- [6] Berry, P E & Fogg, D A B "Optimal Search, Location and Tracking of Surface Maritime Targets by a Constellation of Surveillance Satellites", DSTO Research Report in draft

DISTRIBUTION LIST

A Dynamic Asset Tasking Technique for Integrated Surveillance Operations

Paul E. Berry

AUSTRALIA

DEFENCE ORGANISATION

Task sponsor

Director General Intelligence, Surveillance, Reconnaissance and Electronic Warfare

S&T Program

Chief Defence Scientist	}	shared copy
FAS Science Policy		
AS Science Corporate Management		
Director General Science Policy Development		
Counsellor Defence Science, London (Doc Data Sheet only)		
Counsellor Defence Science, Washington (Doc Data Sheet only)		
Scientific Adviser Joint		
Navy Scientific Adviser (Doc Data Sheet and distribution list only)		
Scientific Adviser - Army (Doc Data Sheet and distribution list only)		
Air Force Scientific Adviser		
Director Trials		

Information Sciences Laboratory

Chief of Intelligence, Surveillance and Reconnaissance Division
Chief of Electronic Warfare and Radar Division
Mr P Amey
Dr A Shaw
Mr D Fogg
Dr J Whitrow
Dr B Haywood
Dr D McDonald
Dr C Pontecorvo
Dr A Pincombe
Dr P Hew
Author(s): Dr PE Berry

DSTO Library and Archives

Library Edinburgh 2 copies
Australian Archives

Capability Systems Staff

Director General Maritime Development (Doc Data Sheet only)
Director General Land Development

Director General Aerospace Development (Doc Data Sheet only)

Knowledge Staff

Director General Command, Control, Communications and Computers (DGC4)
(Doc Data Sheet only)

Navy

SO (SCIENCE), COMAUSNAVSURFGRP, NSW (Doc Data Sheet and
distribution list only)

Army

ABCA National Standardisation Officer, Land Warfare Development Sector,
Puckapunyal (4 copies)
SO (Science), Deployable Joint Force Headquarters (DJFHQ) (L), Enoggera QLD
(Doc Data Sheet only)

Intelligence Program

DGSTA Defence Intelligence Organisation
Manager, Information Centre, Defence Intelligence Organisation

Defence Libraries

Library Manager, DLS-Canberra
Library Manager, DLS - Sydney West (Doc Data Sheet Only)

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy
Library
Head of Aerospace and Mechanical Engineering
Serials Section (M list), Deakin University Library, Geelong, VIC
Hargrave Library, Monash University (Doc Data Sheet only)
Librarian, Flinders University

OTHER ORGANISATIONS

National Library of Australia
NASA (Canberra)
State Library of South Australia

OUTSIDE AUSTRALIA

INTERNATIONAL DEFENCE INFORMATION CENTRES

US Defense Technical Information Center, 2 copies
UK Defence Research Information Centre, 2 copies
Canada Defence Scientific Information Service, 1 copy
NZ Defence Information Centre, 1 copy

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK
Library - Exchange Desk, National Institute of Standards and Technology, US

SPARES (5 copies)

Total number of copies: 47

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA					
				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
2. TITLE A Dynamic Asset Tasking Technique for Integrated Surveillance Operations			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) Paul E Berry			5. CORPORATE AUTHOR Information Sciences Laboratory PO Box 1500 Edinburgh South Australia 5111 Australia		
6a. DSTO NUMBER DSTO-RR-0246		6b. AR NUMBER AR-012-480		7. DOCUMENT DATE September 2002	
8. FILE NUMBER E 9505-23-97		9. TASK NUMBER JTW 99/005		10. TASK SPONSOR JTW 99/005	
				11. NO. OF PAGES 16	
				12. NO. OF REFERENCES 6	
13. DOWNGRADING/DELIMITING INSTRUCTIONS				14. RELEASE AUTHORITY Chief, Intelligence, Surveillance and Reconnaissance Division	
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for Public Release.</i>					
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111					
16. DELIBERATE ANNOUNCEMENT No Limitations					
17. CITATION IN OTHER DOCUMENTS Yes					
18. DEFTEST DESCRIPTORS Surveillance, Defence Planning, Operational Effectiveness, Systems Integration					
19. ABSTRACT This report is concerned with the development of a surveillance dynamic tasking tool whose purpose is to demonstrate benefits in terms of enhanced operational effectiveness through the coordinated deployment and control of a suite of surveillance assets. The underlying theory is developed for mathematically representing surveillance information, and the problem of optimising a search operation so as to maximise the information collected stated formally. A technique for solving the problem is proposed using evolutionary programming and a simplified version of the tool has been implemented so as to control sensors and platforms undertaking a simulated search operation in order to demonstrate the feasibility of the approach.					

